

## Repetition is easy: Why repeated referents have reduced prominence

TUAN Q. LAM AND DUANE G. WATSON  
University of Illinois at Urbana-Champaign, Urbana, Illinois

The repetition and the predictability of a word in a conversation are two factors that are believed to affect whether it is emphasized: Predictable, repeated words are less acoustically prominent than unpredictable, new words. However, because predictability and repetition are correlated, it is unclear whether speakers lengthen unpredictable words to facilitate comprehension or whether this lengthening is the result of difficulties in accessing a new (nonrepeated) lexical item. In this study, we investigated the relationship between acoustic prominence, repetition, and predictability in a description task. In Experiment 1, we found that repeated referents are produced with reduced prominence, even when these referents are unexpected. In Experiment 2, we found that predictability and repetition both have independent effects on duration and intensity. However, word duration was primarily determined by repetition, and intensity was primarily determined by predictability. The results are most consistent with an account in which multiple cognitive factors influence the acoustic prominence of a word.

In conversation, certain words stand out more than others. Acoustically, these words are prominent because they are produced with greater intensity, a higher fundamental frequency ( $f_0$ ), or longer duration than expected. Traditionally, prosodic prominence has been described in two different ways. In one tradition, *prominence* is defined by a linguistic construct called a *pitch accent*, which occurs on words that are new or focused. Pitch accents are typically marked with a change in  $f_0$ , and different types of pitch accents play different roles in the discourse (e.g., Pierrehumbert, 1980; Pierrehumbert & Hirschberg, 1990). Prominence has also been described in terms of its acoustic-phonetic form: Prominence correlates with increases in fundamental frequency ( $f_0$ ), duration, intensity, and intelligibility (Bard et al., 2000; Bell, Brenier, Gregory, Girand, & Jurafsky, 2009; Bell et al., 2003; Fowler & Housum, 1987; Jurafsky, Bell, Gregory, & Raymond, 2001; Watson, Arnold, & Tanenhaus, 2008). These different correlates of prominence often covary; however, they do not perfectly co-occur (e.g., Bard et al., 2000). Within this tradition, variation in the acoustic form has typically been discussed in terms of the word or syllable's predictability or in terms of the psychological factors that play a role in producing or understanding the word. In the present study, we investigated which factors drive apparent effects of predictability and discourse on acoustic prominence and examined whether effects of redundancy on acoustic prominence are a result of speakers' altering their speech to maintain a uniform information profile or whether these effects are linked to language production retrieval and preparation processes.

A number of factors have been shown to correlate with prominence, such as repetition (Aylett & Turk, 2004; Bard et al., 2000; Bard & Aylett, 1999; Bell et al., 2009; Fowler & Housum, 1987; Pluymaekers, Ernestus, & Baayen, 2005a), frequency (Fosler-Lussier & Morgan, 1999; Gregory, Raymond, Bell, Fosler-Lussier, & Jurafsky, 1999; Jurafsky et al., 2001; Pluymaekers, Ernestus, & Baayen, 2005b), and transitional probability (Bell et al., 2009; Jurafsky et al., 2001; Kidd & Jaeger, 2008). For example, Fowler and Housum found that previously mentioned words in a corpus of recorded speech are shorter and less intelligible than words that have not been previously mentioned. Similarly, in recorded speech generated from a referential communication task, Bard and Aylett found that repeated words are less intelligible to listeners than are nonrepeated words, and other work has shown that listeners interpret prominence as a cue to new information in online sentence processing (e.g., Dahan, Tanenhaus, & Chambers, 2002). Lexical frequency is also linked with prominence (Zipf, 1929). High-frequency words are produced with shorter durations than low-frequency words (Fosler-Lussier & Morgan, 1999; Gregory et al., 1999; Jurafsky et al., 2001; Pluymaekers et al., 2005b). Lexical frequency also affects affix duration. When affixes are attached to infrequent words, the affixes are longer than when they are attached to more frequent words (Pluymaekers et al., 2005b). Finally, transitional probability can also affect prominence (Gregory et al., 1999; Jurafsky et al., 2001; Kidd & Jaeger, 2008; Pluymaekers et al., 2005a). When the transitional probability of a word is high, the acoustic realization of the

---

T. Q. Lam, tuanlam2@uiuc.edu

---

word is reduced relative to when the transitional probability is low (Jurafsky et al., 2001).

### Theories of Acoustic Prominence

Although it is clear that repetition, frequency, transitional probability, and predictability in general all affect prosodic prominence, it is less clear why these effects exist. One proposal is that prominence differences are the result of speakers' optimizing the acoustic signal for comprehension (Aylett & Turk, 2004; Fowler & Housum, 1987; Frank & Jaeger, 2008; Lieberman, 1963). If words are new, infrequent, and less predictable, they may be difficult to identify in running speech. As a result, speakers may articulate these words more clearly to facilitate processing by the listener. Repeated, frequent, and predictable words are readily identifiable, so there is less need to articulate these words carefully (Fowler & Housum, 1987).

More recently, some accounts of prominence have appealed to information-theoretic principles to explain differences in prominence across words. One example is Aylett and Turk's (2004) smooth signal redundancy hypothesis. According to this account, effects of repetition, frequency, and predictability on duration can be reduced to one thing: language redundancy (Aylett & Turk, 2004). *Language redundancy* is the predictability of a syllable, word, or syntactic structure in a linguistic context. It is modulated by a number of factors, including lexical frequency, syntax, and pragmatics. In this framework, effects of repetition can also be linked to redundancy. Because speakers are more likely to refer to previously mentioned referents than to new referents (Arnold, 1998), the repeated mention of a word is more predictable than the mention of a new word.

According to the smooth signal redundancy hypothesis, speakers attempt to produce a signal in which the amount of redundancy remains relatively constant throughout production. Aylett and Turk (2004) proposed that prosodic prominence's primary role is to smooth the information profile of a word. This is accomplished by reduction of syllable duration when a word is redundant: Expected words are produced with shorter durations, and unexpected words are produced with longer durations, so that the amount of information conveyed is evenly distributed over time. Thus, the smooth signal redundancy hypothesis can explain the effects of repetition, frequency, and contextual predictability on prominence reduction.

Although information-theoretical accounts argue that reduction is for the benefit of the listener, this does not entail that speakers explicitly model individual listeners such that specific knowledge is linked to a particular listener. Rather, speakers may model a generic listener (E. A. Isaacs & Clark, 1987) whom they expect will have a certain set of beliefs or expectations. Taken to the extreme, speakers may model predictability for any listener in general (Brown & Dell, 1987). Proponents of this approach argue that it is a model of communication at the computational level, rather than at the algorithmic level, in a Marr-like (1982) framework, and they are agnostic as to how it is implemented psychologically (e.g., Aylett & Turk, 2004; Frank & Jaeger, 2008; Jaeger, 2010). Critically, whether a

specific listener is modeled or a more generic listener is modeled, the approach at its core depends on smoothing the signal for a listener.

An alternative view is that the link between predictability and acoustic prominence is the result of speaker-internal production processes and not a means by which the speaker facilitates processing for the listener (Bard et al., 2000; Bell et al., 2009). Most researchers agree that speech production is a multistep process beginning with message formulation, followed by grammatical encoding, and finishing with phonological encoding (e.g., Bock & Levelt, 1994). During *message formulation*, speakers formulate the semantic meaning of what they are going to say. During *grammatical encoding*, speakers select the appropriate lexical items to convey their messages, compute word order, and include lexical items (like function words) that satisfy the constraints of the grammar of the language that they speak. After grammatical encoding, speakers must encode the linguistic material into a phonological representation.

Proponents of a speaker-internal account of prominence argue that prominence is linked to the amount of activation associated with a word in lexical retrieval (e.g., Bell et al., 2009). The speed at which lexical items are retrieved in the course of language production is regulated by factors like word frequency, repetition, and contextual predictability: Frequent, repeated, and predictable words are retrieved more quickly than infrequent, nonrepeated, and less predictable words (Griffin & Bock, 1998; Jescheniak & Levelt, 1994). Bell et al. (2009) proposed that the speed of lexical retrieval is linked to articulatory planning, such that words that are retrieved quickly are articulated more quickly, whereas words that are retrieved slowly are produced slowly. Bell et al. (2009) argued that this coordination between the speed of lexical retrieval and articulatory planning is a strategy used by the production system to maximize fluent speech.

Thus, under a lexical retrieval account, repeated, predictable, and frequent words are reduced, because they are retrieved quickly, whereas nonrepeated, unpredictable, and infrequent words are more prominent because they are retrieved more slowly. For example, when an entity is new to discourse, it will typically be less expected (Arnold, 1998). As a result, there will be very little activation of the lexical item representing that entity, and retrieval will be slower than if the word was initially more activated (Jescheniak & Levelt, 1994; Wingfield, 1968). This should result in longer articulation of the word. However, a given entity will have been previously activated and may be less difficult to retrieve than when it was first mentioned. This could occur either because of maintenance or through the slow decay of activation (Dell, 1990). Ultimately, this would lead to acoustic reduction of the word.

Although Bell et al.'s (2009) proposal centered on lexical access as the primary determinant of word duration, in principle, facilitation of processing at other levels of production could also affect word duration. If the linguistic message is easily formulated because of its frequency, the referential context, or because it is repeated, one might expect reduction. Similarly, if the phonological form of a

word is highly activated, this might lead to reducing the phonological form as well. At the heart of all of these proposals is a desire to maximize fluent speech through feed-forward mechanisms from earlier stages of production to the articulatory planning stage. We return to the locus of these potential effects in the General Discussion section.

Finally, a third potential account of acoustic prominence is the multiple-source view (Watson, 2010). Rather than assuming that acoustic prominence has a single source, under the multiple-source view, the acoustic realization of a word is the product of many factors, including difficulty in speech production and marking information for a listener. These different factors may differentially affect varying aspects of the acoustic signal. This approach differs from that of the information-theoretic and the lexical-access accounts, which make specific predictions about reduction (i.e., the shortening of the duration of a word). Under the multiple-source account, a word's fundamental frequency, duration, and intensity may be affected by different factors in different ways. Evidence for this view comes from both the production and the comprehension literature. Watson et al. (2008) found that difficult moves in games of Tic Tac Toe are produced with longer duration than are moves that are easy or predictable. Moves that are important to the game, like a winning move or blocking a winning move, are produced with greater intensity than those that are not. In comprehension, some acoustic cues are used preferentially over other cues in determining linguistic structure, suggesting that these differing cues might have differing underlying sources (A. M. Isaacs & Watson, 2009, 2010). For example, listeners use the  $f_0$  slope over a word rather than raw duration in detecting prominence that is linked to discourse status (A. M. Isaacs & Watson, 2010). A. M. Isaacs and Watson (2009) found that intensity and not duration contributes to metalinguistic judgments of acoustic prominence, even though both correlate with prominence in production. Therefore, it is possible that acoustic prominence might be the result of both speaker-internal production mechanisms and facilitating comprehension, and these factors may affect the acoustic signal in different ways.

Although all three accounts predict that predictability, lexical frequency, and repetition will influence acoustic prominence, this study is focused on repetition. This is because a lexical-retrieval account and an information-theoretic account can potentially make differing predictions about whether repeated words will be reduced. Under a lexical-retrieval account, repetition is critical for reduction. Words that have been produced before are reduced because they are easy to repeat. In contrast, under the information-theoretic approach, reduction is driven primarily by how expected a word is. One way to test whether reduction is the result of information-theoretic principles or the result of processes related to lexical retrieval and planning is to test whether effects of repetition and predictability are independent. If prominence is the product of smoothing the information profile of a word for the listener, unexpected words should be lengthened, even if they have been repeated. If prominence is partly the result of speaker-internal lexical-retrieval processes,

speakers should reduce previously mentioned words and lengthen new words, independent of whether the word is expected. Residual activation stemming from the previous production should cause reduction, even if the word is unexpected. Finally, the multiple-source account allows for the possibility that both theories may play some role in acoustic prominence, possibly in different ways.

Although we have discussed the acoustic correlates of prominence very generally, note that the information-theoretic account and lexical-retrieval account make predictions about duration in particular. Under information-theoretic accounts, changing word duration is the means by which speakers alter the word's information profile. Similarly, under lexical-retrieval theories, the difficulty of producing a word affects the word's length. In the experiments reported below, differences in duration will be used to adjudicate between these theories, although we also measure  $f_0$  and intensity in order to determine whether they too are linked to repetition and predictability. The latter is critical for testing the multiple-source account, because it is possible that predictability and repetition might both have effects on acoustic prominence, but they may occur along different acoustic dimensions.

Previous studies of repetition and predictability have relied primarily on corpus data (e.g., Aylett & Turk, 2004; Bell et al., 2009; Jurafsky et al., 2001). However, in natural speech, repetition and predictability are highly correlated: Repeated words are more predictable than nonrepeated words (Arnold, 1998). Therefore, it is difficult to know whether effects of repetition and predictability are the result of similar cognitive processes. In Experiments 1A and 1B, we address this question by altering the correlation between predictability and repetition that exists in natural speech. In a picture description task, contexts were created in which repeating a word was unexpected, and producing a new word was expected (Experiment 1A). Contexts were also created in which repeating a word and producing a new word were equally expected, in order to elicit responses for comparison with those from Experiment 1A (Experiment 1B). If repeated reference causes reduction even when that target is less predictable, this would provide support for theories that attribute reduction to factors in lexical retrieval. If repeated, less predictable target words are produced with longer durations than are nonrepeated, predictable words, this would support a redundancy-avoidance account: Predictable words are reduced to facilitate robust communication for the listener. If both repetition and predictability play a role, this would provide evidence for the multiple-source account.

## EXPERIMENTS 1A AND 1B

### Method

**Participants.** Sixty-three undergraduate students from the University of Illinois at Urbana-Champaign participated in this experiment to earn credit in a psychology course (32 in Experiment 1A and 31 in Experiment 1B). All of the participants were native speakers of American English. Five participants had to be excluded from the analysis: 1 participant failed to produce the second utterance on repeated trials. Another 2 participants were excluded because of a recording error. The remaining 2 participants used pronouns

on repeated mention trials, which made it impossible to compare prominence on repeated and nonrepeated trials.

**Materials.** The participants' task was to describe events on a computer screen to a confederate. Two pictures appeared on the participant's (the director) and the confederate's (the matcher) screen for each trial. On a given trial, one of the objects would shrink, and then one of the objects would flash.

The stimuli were taken from a set of 12 images from Rossion and Pourtois (2001). These images were a colorized version of images originally created by Snodgrass and Vanderwart (1980). The images were used to generate six pairs of images. The images were paired to avoid semantic and phonetic relatedness. The image pairs were presented side by side in the center of the screen (see Figure 1). Each image pair appeared 18 times during the experiment, for a total of 108 trials. Which image appeared on the left or right was counterbalanced, such that each image appeared on each side an equal number of times. The items were randomized within sets of 12 trials such that in each set, each image pair appeared twice to counterbalance the image location.

Repetition of one of the objects in the task and the likelihood of its being mentioned were both manipulated. Repetition was manipulated by varying whether the same object engaged in a shrinking and a flashing event. On repeated-mention trials, one of the images shrank, and then the same image flashed. On nonrepeated-mention trials, one of the images shrank, and then the other image flashed. An example of a trial from the nonrepeated condition is shown in Figure 1.

Typical utterances are presented in Examples 1 and 2.

- (1) Repeated noun  
The axe is shrinking . . . The axe is flashing.
- (2) Nonrepeated noun  
The penguin is shrinking . . . The axe is flashing.

In order to manipulate predictability, a training block, followed by a test block, was used to manipulate speaker expectations about what object on the screen would flash. The training block consisted of 96 trials, and the test block consisted of 12 trials. In the training block of Experiment 1A, the repeated-mention trials were much less predictable than the nonrepeated-mention trials. In the training block, only 6 of the trials were repeated-mention trials. The order of the conditions was pseudorandomly permuted, such that no repeated-mention trial occurred within 5 trials of another repeated-mention trial. Moreover, the first 8 trials were all nonrepeated-mention trials. In the training block of Experiment 1B, repeated-mention trials and nonrepeated-mention trials were equally likely, and the order of the conditions was randomly permuted. Each training set was random-

ized such that every participant was presented with a different item order during the training block.

The purpose of the test block was to determine what effects repetition and the expectations established in the training block had on speaker productions. The test block consisted of 12 trials. The transition between the training and test blocks was not marked, so the participants were unaware of the transition. During the test block, the participants were exposed to each pair of images twice, once in the repeated condition and once in the nonrepeated condition. As with the training block, the order of items during the test block was randomized. No pair of items appeared twice in a row during the test block. In Experiment 1A, there were 30 unique lists, such that each participant was exposed to a unique order of trials during the test block. The order in which items were presented in the test block was matched across Experiments 1A and 1B.

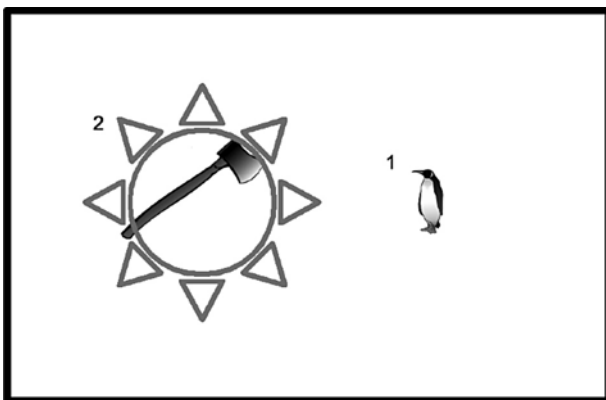
The experiments were programmed using MATLAB with the Psychophysics Toolbox Version 2.54 installed. Participant utterances were recorded at a frequency of 44 kHz.

**Procedure.** Before the experiment, the participants were shown a video in which two research assistants were presented completing the task. This video was used both to instruct the participants on how to complete the task and to prime the participants with the construction *the [Noun 1] is shrinking . . . The [Noun 2] is flashing*. Then the participants were told that they would play the role of the director, and the research assistant would play the role of the matcher. The director sat at a computer facing away from the matcher's computer. The pairs completed six practice trials before beginning the actual experiment. The images in the practice trials were the same images as those used in the experimental trials, and each pair was used only once. All of the practice trials were nonrepeated-mention trials.

At the beginning of each trial, two images appeared. After 500 msec, one of the images shrank. Two seconds after the shrinking event, either the same image flashed or the other image flashed. The participants were instructed to describe each event as soon as they knew what was happening. The matcher, meanwhile, clicked on one of four buttons to make her screen match the director's screen. These buttons corresponded to the shrinking event and the flashing event for each image. The trial ended when the matcher notified the director that she was finished matching the second event. Occasionally, the speaker accidentally misnamed an object. On those trials, the matcher provided feedback that she did not have the object, thereby prompting the director to correct his or her utterance. Otherwise, there was no explicit feedback aside from confirmation of completion of the trial. The pair completed 108 trials, of which the last 12 were recorded and labeled using Praat, a speech analysis program developed by Boersma and Weenink (2007). The target words during the first and second utterances were analyzed for mean  $f_0$  over the word, the maximum  $f_0$  excursion over the word, the minimum  $f_0$  over the word, word duration, and mean intensity. We also computed the proportion of the total utterance duration due to the target word in each condition (hereafter called *target proportion*).

## Results

The data were analyzed using a linear mixed-effects regression model with subject and item as random intercepts and slopes using the lmer function in the lme4 package in R (Baayen, 2008). Like ANOVA, this method accounts for the variance due to subjects and items; however, this method can account for variance of multiple random factors simultaneously (Baayen, 2008). Model comparisons were conducted using a likelihood ratio test to find the best-fit random slopes and intercepts models. Random slopes did not significantly increase model fit for any reported model and are therefore not reported. The reported  $p$  values were obtained from Markov chain Monte Carlo sampling using the R language package (Baayen, 2008).



**Figure 1.** Illustration of a typical trial from the nonrepeated condition in Experiment 1. The numbers indicate the order of events. The sun indicates the flashing event.



**Table 1**  
**Experiments 1A and 1B Noun 2 Summary**

Metric	Experiment 1A				Experiment 1B			
	Nonrepeated		Repeated		Nonrepeated		Repeated	
	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>
Duration (msec)	393	13.8	376	14.3	356	10.3	339	10.7
Proportion	.424	.008	.413	.008	.420	.007	.406	.007
Intensity (dB)	78.7	1.14	78.6	1.08	77.6	0.97	77.0	1.00
Average <i>f</i> 0 (Hz)	167	9.02	165	9.14	156	8.94	155	8.19
<i>f</i> 0 maximum (Hz)	196	11.3	194	11.2	179	10.7	181	10.7
<i>f</i> 0 minimum (Hz)	141	8.07	140	8.23	136	7.36	135	7.19

Note—In Experiment 1A, repeated nouns had low predictability (6.25% of trials), whereas nonrepeated nouns had high predictability (93.75% of trials). In Experiment 1B, repeated and nonrepeated nouns were equally likely.

The production of the second target word across conditions was compared. We also compared the production of the first target word in a trial (Noun 1) with the production of the second target word (Noun 2) for trials in the repeated condition. All predictor variables were contrast coded, and as a result of our balanced design, the predictors were centered around the mean.

In Experiment 1A, there were effects of repetition but no effects of predictability. Nonrepeated Noun 2s were produced with greater duration ( $t = 3.421, p < .001, \beta = 17.0, SE = 4.96$ ) than were repeated Noun 2s; however, there was no significant difference for intensity ( $t < 1$ ). There was also no significant difference in *f*0 across conditions for Noun 2s. The target proportion in the nonrepeated condition was also significantly greater than the target proportion in the repeated condition ( $t = 2.969, p < .001, \beta = 0.0112, SE = 0.00377$ ). The means for Noun 2 across repetition conditions are presented in Table 1. In the repeated trials, Noun 1 was produced with greater duration ( $t = 3.78, p < .001, \beta = 20.9, SE = 5.54$ ) and intensity ( $t = 4.98, p < .0001, \beta = 1.08, SE = 0.216$ ) than was Noun 2. Noun 1 was also produced with a higher maximum *f*0 ( $t = 2.98, p < .01, \beta = 13.2, SE = 4.434$ ), a higher minimum *f*0 ( $t = 4.91, p < .001, \beta = 18.7, SE = 3.82$ ), and a higher mean *f*0 ( $t = 6.13, p < .001, \beta = 14.6, SE = 2.38$ ) than Noun 2. The means for Noun 1 and Noun 2 in the repeated condition are presented in Table 2.

In Experiment 1B, nonrepeated Noun 2s were produced with greater duration ( $t = 3.62, p < .001, \beta = 16.9, SE =$

4.68) and intensity ( $t = 2.83, p < .01, \beta = 0.645, SE = 0.228$ ) than were repeated Noun 2s. The target proportion was also significantly larger in the nonrepeated condition ( $t = 3.89, p < .0001, \beta = 0.0139, SE = 0.00356$ ). As in Experiment 1A, there was no significant difference in *f*0 across conditions for Noun 2. In the repeated trials, Noun 1 was produced with greater duration ( $t = 2.76, p < .01, \beta = 12.7, SE = 4.62$ ) and intensity ( $t = 5.06, p < .0001, \beta = 1.52, SE = 0.231$ ) than was Noun 2. Noun 1 was also produced with a higher average *f*0 ( $t = 4.07, p < .05, \beta = 10.7, SE = 2.64$ ) and a higher minimum *f*0 ( $t = 3.05, p < .01, \beta = 10.8, SE = 3.53$ ). Maximum *f*0 did not differ significantly across conditions.

**Discussion**

In both experiments, repeated nouns were less prominent than nonrepeated nouns, providing support for a lexical-retrieval account of prosodic prominence. According to these accounts, repeated words should be reduced, because they have been previously activated and are therefore easier to retrieve for a subsequent production. This was also true of the overall utterance duration and the target proportion. Surprisingly, there were no differences among Noun 2s in *f*0. In the repeated conditions, *f*0 was higher for the first noun than for the second noun, although this may have been due to the declination in pitch that typically occurs over a set of related utterances. It is possible that the descriptive nature of the task led the participants to vary their pitch less than they would have in a more interactive setting.

**Table 2**  
**Means of Noun 1 and Noun 2 in the Repeated Condition in Experiments 1A and 1B**

Metric	Experiment 1A				Experiment 1B			
	Noun 1		Noun 2		Noun 1		Noun 2	
	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>
Duration (msec)	397	14.0	376	13.8	352	10.6	339	10.7
Intensity (dB)	79.7	1.10	78.6	1.08	78.5	0.95	77.0	1.00
Average <i>f</i> 0 (Hz)	180	10.1	165	9.14	165	9.60	155	8.19
<i>f</i> 0 maximum (Hz)	207	12.0	194	11.2	188	10.7	181	10.7
<i>f</i> 0 minimum (Hz)	159	8.98	140	8.23	145	8.25	135	7.19

Note—In Experiment 1A, repeated Noun 2s had low predictability (6.25% of trials). In Experiment 1B, repeated Noun 2s had relatively higher predictability (50% of trials).

These results are less consistent with information-theoretic approaches. According to information-theory-based accounts, the predictability manipulation in Experiment 1A should have led to reduced duration of nonrepeated, expected nouns and lengthening of repeated, unexpected nouns. In fact, the reverse occurred.

Note, however, that although the duration results support the theory that prominence is a result of lexical retrieval production processes, there were some differences in intensity across experiments. A post hoc test for a condition  $\times$  experiment interaction was conducted on the Noun 2 intensity data, which yielded a marginally significant interaction ( $t = 1.77, p = .07, \beta = 0.561, SE = 0.317$ ). There were no differences in intensity in Experiment 1A between the repeated, unexpected condition and the nonrepeated, expected condition, but in Experiment 1B, nonrepeated Noun 2s were produced with greater intensity than repeated Noun 2s. This pattern of results suggests that the lack of intensity differences observed in Experiment 1A may have been due to the predictability manipulation. This marginal interaction suggests that predictability and repetition may in fact be two separate factors that influence the production of prominence, a possibility that is most consistent with the multiple-source view of prominence outlined in the introduction.

Predictability and repetition might both have had effects on intensity, but effects of predictability were not detectable, because predictability and repetition were not independently manipulated. In the contexts of these experiments, it is difficult to determine whether this occurred, because in both instances, predictability and repetition were negatively correlated. We address this issue in Experiment 2.

## EXPERIMENT 2

The goal of Experiment 2 was to test whether predictability and repetition have separate, independent effects on the acoustic realization of a word. A shortcoming of Experiment 1A was that predictability and repetition were negatively correlated, so the effects of one factor might have obscured the effects of the other.

In Experiment 2, we altered the task used in Experiments 1A and 1B so that predictability and repetition could be independently manipulated. The participants were presented with an array of 12 images. As in Experiments 1A and 1B, one image shrank, and then another image flashed. However, the second event was preceded by a probabilistic cue as to which object would flash. In 92% of the trials, a circle appeared around the image that flashed. At the beginning of the experiment, the participants were told that the circle would usually but not always indicate which object would flash. This made it possible to independently manipulate whether a given word was repeated across events and whether it was expected by the speaker. An information theoretic account predicts that expected words will be produced with shorter durations than will unexpected words. This account also predicts that there should be no differences between repeated and nonrepeated words, be-

cause the reliability of the cue was the same in the repeated-mention and the nonrepeated-mention trials. In contrast, a lexical-retrieval account predicts that repetition, rather than predictability, should affect reduction. This account predicts that repeated words should be produced with less prominence than nonrepeated words. Finally, the multiple-source account predicts that both repetition and predictability affect prominence and that these effects may be realized in different ways acoustically.

## Method

**Participants.** Forty-five undergraduate students from the University of Illinois at Urbana-Champaign participated in this study in exchange for course credit. All of the participants were native speakers of American English. Four participants were excluded because of recording errors. One participant was excluded for failing to follow the instructions.

**Materials.** As in Experiments 1A and 1B, images were taken from a set of colored images by Rossion and Pourtois (2001). Ninety-six images from this set were used, of which 72 images were targets in critical trials. For each trial, a  $3 \times 4$  array of images was displayed on a computer screen using MATLAB with the Psychophysics Toolbox Version 3.0 installed.

As in Experiments 1A and 1B, there were two events on each trial. On every trial, one of the images shrank, and one of the images flashed. Predictability and repetition were manipulated in a  $2 \times 2$  factorial design. Predictability was manipulated by circling a potential target for the second utterance immediately after the shrinking event but before the flashing event. On predictable trials, the circled image flashed. On unpredictable trials, the circled image did not flash. Repetition of the target word was also manipulated: Either the same object shrank and flashed or different objects shrank and flashed. Predictability and repetition were crossed in a  $2 \times 2$  design yielding four conditions: repeated expected, repeated unexpected, nonrepeated expected, and nonrepeated unexpected.

At the beginning of a trial, 12 images appeared on the screen. After 1 sec, one of the images shrank. Then, after 1 sec, one of the images was circled. The circle remained on the screen for 500 msec and then disappeared. After another 500 msec, one of the images flashed. Which images shrank and/or were circled depended on the condition for the trial. In the repeated-mention trials, the image that shrank was also the image that flashed. In the nonrepeated-mention trials, the image that shrank was not the image that flashed. In the expected trials, the image that was circled was the image that eventually flashed. In the unexpected trials, the image that was circled was not the image that eventually flashed. An illustration of a trial from the nonrepeated unexpected condition is shown in Figure 2.

There were 6 critical trials for each condition, for a total of 24 critical trials. There were 120 filler trials, which were all trials in which the targets were nonrepeated and predictable. These filler trials were used to reinforce the predictability manipulation. Overall, the cue was reliable on 132 of 144 trials, or roughly 92% of the time. The order of trials was pseudorandomly permuted, such that no two critical trials appeared in succession. On the critical trials, the shrinking image, the circled image, and the flashing image were all novel as targets but may have appeared previously as filler images. After the critical trial in which these images were used as targets, they could appear again later as targets in filler trials. Because of the potential for order effects, two pseudorandomized target lists were used. In each list, the critical items were counterbalanced using a Latin square, resulting in eight lists in total. The target on each trial was counterbalanced so that it appeared in all 12 locations an equal number of times in both critical trials and filler trials.

**Procedure.** Before beginning the experiment, the participants were shown a video that demonstrated the task. This video was used to prime the participants to use the construction *the* [Noun 1] *is shrinking* . . . *The* [Noun 2] *is flashing*. The video was also used




Figure 2. Example of a typical trial from the nonrepeated expected condition in Experiment 2. The numbers indicate the order of events in the trial. The sun indicates the flashing event.

to inform the participants about the probabilistic cue to the flashing event. The participants were told that the circled image was frequently the image that flashed. After watching the video, the participants completed eight practice trials. Five of the trials were generated from the nonrepeated predictable condition. The other three trials were from the remaining conditions (repeated expected, repeated unexpected, and nonrepeated unexpected). These last three conditions were included so that the participants would not be surprised when they encountered them in the actual experiment.

Following the practice trials, the participants immediately began the experiment. Unlike in Experiments 1A and 1B, in which the participants addressed a matcher, in Experiment 2, the speaker was alone in the room, and progress through the experiment was self-paced. At the beginning of a trial, one of the images shrank. The participants then described this event. Then one of the images was circled. Finally, one of the images flashed and the participants described this event. After describing the flashing event, the participants pressed a key to begin the next trial. The participants completed 144 continuous trials, of which 24 critical trials were recorded. The target word was the production of the flashing noun. The mean *f*<sub>0</sub>, maximum *f*<sub>0</sub>, minimum *f*<sub>0</sub>, intensity, and duration of target words were measured. We also computed the proportion of the total duration of the utterance occupied by the target word in each condition.

**Results**

Two targets (screw and refrigerator) were removed from the analysis because of inconsistency in naming them across participants. Of the remaining targets, 47 trials were removed from the analysis because of errors in

naming the targets. This led to a loss of 5.34% of the total trials.

Means are presented in Table 3. The data were analyzed using linear mixed-effects regression with participant and item as random effects. All predictor variables were contrast coded and centered. As in Experiments 1A and 1B, model comparisons were conducted using a likelihood ratio test to find the best-fit random slopes and intercepts models. Again, random slopes did not significantly increase model fit for any reported model and are therefore not reported. All of the reported models include random intercepts.

Both repetition and expectedness were reliable predictors of intensity, and the two factors together were better predictors of intensity than either factor alone. Nonrepeated words had greater intensity than repeated words ( $t = 2.73, p < .01, \beta = 0.4802, SE = 0.176$ ), and unexpected words had greater intensity than expected words ( $t = 3.77, p < .001, \beta = 0.6617, SE = 0.176$ ). The relative size of the regression coefficients suggests that predictability had a larger effect than repetition on intensity.

Only repetition was a reliable predictor of raw noun duration. Nonrepeated words were longer than repeated words ( $t = 6.58, p < .0001, \beta = 30.0, SE = 4.55$ ). Predictability did not reliably predict raw duration.

Table 3  
Experiment 2 Noun 2 Summary

Metric	Nonrepeated				Repeated			
	Expected		Unexpected		Expected		Unexpected	
	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>	<i>M</i>	<i>SEM</i>
Duration (msec)	446	13.5	458	11.9	421	10.9	428	12.1
Noun proportion	.401	.006	.420	.006	.398	.005	.404	.006
Intensity (dB)	58.4	0.81	59.2	0.88	58.2	0.82	58.7	0.91
Average <i>f</i> <sub>0</sub> (Hz)	170	7.49	177	7.62	173	6.84	168	6.70
<i>f</i> <sub>0</sub> maximum (Hz)	206	9.28	220	9.54	213	10.0	198	8.87
<i>f</i> <sub>0</sub> minimum (Hz)	146	6.51	149	6.71	145	5.84	144	6.07

There were effects of both repetition and predictability on target proportion. Nonrepeated words had a greater target proportion than repeated words ( $t = 5.37$ ,  $p < .0001$ ,  $\beta = 0.0145$ ,  $SE = 0.00270$ ), and unexpected words had a greater target proportion than expected nouns ( $t = 2.70$ ,  $p < .01$ ,  $\beta = 0.0073$ ,  $SE = 0.00270$ ). For target proportion, the relative size of the regression coefficients suggests that repetition has a larger effect than predictability.

For measures of maximum  $f_0$  and average  $f_0$ , there was a significant interaction between repetition and expectedness ( $t = 2.954$ ,  $p < .01$ ,  $\beta = 25.1$ ,  $SE = 8.51$ , and  $t = 2.362$ ,  $p < .05$ ,  $\beta = 8.57$ ,  $SE = 3.65$ , respectively). Post hoc paired  $t$  tests revealed that nonrepeated nouns had a higher average  $f_0$  than did repeated nouns in the unexpected condition [ $t(39) = 3.367$ ,  $p < .001$ ]; expected nouns showed no difference across repetition conditions [ $t(39) = 0.088$ ,  $p = .93$ ].

### Discussion

The results from Experiment 2 suggest that repetition and predictability play independent roles in the production of prosodic prominence. Both factors affected the duration and intensity of the target word. However, predictability was the stronger predictor of intensity, whereas repetition was the stronger predictor of duration. These data are most consistent with the multiple-source account.

The data from Experiment 2 suggest that the effects of repetition obscured effects of predictability in Experiment 1A, particularly for intensity. Recall that in Experiment 1A, these two factors were placed in opposition to one another, making it difficult to know whether each factor contributed independently to the acoustics of the target word.

Interestingly, there was a reliable interaction between predictability and repetition with respect to  $f_0$ , such that unexpected words were produced with higher  $f_0$  in the nonrepeated condition than in the repeated condition, but expected words were produced with similar  $f_0$  in the repeated conditions. This suggests that the lack of an effect in Experiments 1A and 1B was not due to the nature of the task. This interaction was not predicted by any of the theories discussed above and suggests that the presence of both factors may be necessary for triggering a higher  $f_0$ . Future work will be needed to investigate why effects on  $f_0$  appear to be qualitatively different from effects on duration and intensity.

Finally, one potential concern is that because there was no overt listener (unlike in Experiments 1A and 1B), these data may not be useful in evaluating information-theoretic approaches. There are two things to note. The first is that, as was discussed in the introduction, in the information-theoretic frameworks that have been proposed, it is not critical that a specific listener be present. Frank and Jaeger (2008) were agnostic as to whether a specific listener is modeled and argued that these information-theoretic principles apply to communication more generally. The second is that despite the absence of a listener, there were clear effects of predictability. The manipulation of predictability was clearly strong enough to elicit differences.

However, these differences appeared primarily in measures of intensity, rather than duration.

### GENERAL DISCUSSION

In the two experiments discussed in this article, both repetition and predictability influenced the duration and intensity of target words. Repetition primarily affected word length, and predictability primarily affected word intensity. These results are most consistent with a multiple-source view of acoustic prominence: The prominence of a word is affected by production factors like the lexical access account and by marking unpredictable information. In English, these two factors appear to affect the acoustic properties of the word in different ways.

In fact, these findings are consistent with previous findings in the literature. Watson et al. (2008) found a dissociation between intensity and duration in games of Tic Tac Toe, depending on the likelihood and importance of a game move. Baker and Bradlow (2009) found that in clear speech, second mentions are more reduced for high-frequency words than for low-frequency words, which is consistent with repetition and frequency effects having different underlying sources.

The duration results cannot be explained by an information-theory account alone. These theories predict that speakers alter the duration of words such that listeners can more readily parse their utterances depending on the redundancy of the elements in the utterance. In particular, this account suggests that speakers should increase the length of less predictable words and reduce the length of predictable words. However, in Experiment 1A, when repeated words were unexpected, they were still reduced, despite being unexpected. In Experiment 2, in which predictability and repetition were independently manipulated, predictability had no effect on target word duration. Although predictability did have an effect in one of the measures of duration (target proportion), it was weaker than the effects of repetition. Therefore, it is clear that predictability alone is not sufficient for accounting for prominence differences in duration.

One potential concern is that the predictability manipulation here differs from the types of linguistic predictability that have typically been discussed in the literature. Linguistic predictability has been claimed to incorporate lexical and syntactic frequency,  $n$ -gram probabilities, and previous mention, all of which are properties of a language that the native speaker must learn through a lifetime of experience. In contrast, the manipulation of predictability in the present study is based on the predictability of events occurring in a task. The representations that underlie this type of task-based predictability may differ from representations that underlie predictability on the basis of stored linguistic experiences. However, stored sources of predictability are necessarily based on input that the speaker received while interacting with his or her environment and language community. Therefore, in principle, there is no reason that the manipulations of predictability in this task should differ from longer term linguistic predictability, except that it is more recent.



Although the manipulation of predictability did not have a significant effect on the raw duration of the target word, this does not mean that speakers do not optimize some aspects of speech for processing by the listener. First, the speakers did lengthen target proportions, a prediction made by information-theoretical accounts. This can be done either by increasing the duration of the target word or by reducing the duration of words around the target word, with the end result being that the target word is perceived as being more prominent. Although this effect was relatively weak compared with the effects of repetition, it may have resulted from some optimization of the signal for the listener. Second, speakers could improve word intelligibility by increasing the intensity of less predictable words. Changes in intensity may be more useful for a listener. Changes in duration can have a number of underlying causes, including marking prosodic boundaries, the production of disfluencies, marking the rhythm of a sentence, and distinguishing between vowels of varying length, to name just a few. These factors may make duration too noisy of a signal to use in detecting unpredictable information. In contrast, intensity may be a less crowded channel on which to convey information about predictability. And of course, the data here also do not rule out the possibility that changes in duration can be explained by information-theoretic accounts of other aspects of linguistic structure.

One question that we did not address in this study is where these effects arise. Previous work suggests that the repetition effect is not realized at the phonological level. Repetitions lead to reduction, but saying a word and then its homonym does not (Fowler, 1988). Homonyms are words that are identical in spelling and sound but have different meanings. Because homonyms are identical in sound, the production process for homonyms should be identical at the level of phonological encoding. This suggests that the repetition effect is situated at a higher level of production than phonological encoding, potentially at the level of lexical selection or message formulation. Fowler also failed to find a repetition effect when speakers produced repeated words that were produced as a list. Although the discrepancy between Fowler's findings and the results here are puzzling, one possible explanation is that word list production does not engage the same production processes, such as message planning, as situated language use. If reduction is linked to ease of processing at higher stages of the production process, one might not expect to see reduction in the production of word lists. At the very least, both of Fowler's findings suggest that the effects of repetition may be driven by production factors earlier in the production process than the level of phonological encoding. In addition, work showing that reduction and decreases in intelligibility can occur even when a word is produced by a different speaker suggests that these effects are not necessarily rooted in phonological encoding (Anderson & Howarth, 2002; Bard et al., 2000).

It is unclear at what level of production the effects of predictability arise. One possibility is that there are feed-forward connections from the message formulation level to the level of articulation that modulate levels of intensity

depending on the predictability of the word. Judgments of predictability could come either from explicitly modeling the expectations of the listener or from evaluating listeners' knowledge on the basis of the speaker's own assessment of predictability (e.g., Brown & Dell, 1987; Horton & Keysar, 1996). This is related to a broader debate in the psycholinguistics literature regarding the extent to which speakers design their utterances for the listener. Clearly, future work in this domain will be needed to determine the exact mechanism that underlies marking unpredictable information with intensity.

More generally, these data suggest that prominence is not a unitary linguistic or psychological construct. Different factors can play a role in whether a word is produced with prominence, and this prominence can be realized in different ways. In contrast to previous work, in which increases have typically been found in duration and intensity co-occurring in natural speech, we found that intensity is more strongly linked to speaker expectation, whereas duration is more strongly linked to repetition. These data suggest that the prominence of a word can potentially have multiple sources (Watson, 2010).

#### AUTHOR NOTE

We thank Allison Potter, Cara Ader, Nicole Nash, Ashley Turk, and Carl Ruark for helping with data collection and transcription. We also thank Jennifer Arnold for assistance in formulating the initial ideas for this project and Gary Dell and the members of the Communication and Language Lab for helpful comments on this project. This project was supported by Grant R01 DC008774 from the National Institutes of Health. Correspondence concerning this article should be addressed to T. Q. Lam, Department of Psychology, University of Illinois at Urbana-Champaign, 603 E. Daniel St., Champaign, IL 61820 (e-mail: tuanlam2@uiuc.edu).

#### REFERENCES

- ANDERSON, A. H., & HOWARTH, B. (2002). Referential form and word duration in video-mediated and face-to-face dialogues. In J. Bos, M. E. Foster, & C. Matheson (Eds.), *Proceedings of the Sixth Workshop on the Semantics and Pragmatics of Dialogue* (pp. 13-20). Edinburgh: University of Edinburgh.
- ARNOLD, J. E. (1998). *Reference form and discourse patterns*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.
- AYLETT, M., & TURK, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language & Speech*, *47*, 31-56. doi:10.1177/00238309040470010201
- BAAYEN, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge: Cambridge University Press.
- BAKER, R. E., & BRADLOW, A. R. (2009). Variability in word duration as a function of probability, speech style, and prosody. *Language & Speech*, *52*, 391-413. doi:10.1177/0023830909336575
- BARD, E. G., ANDERSON, A. H., SOTILLO, C., AYLETT, M., DOHERTY-SNEDDON, G., & NEWLANDS, A. (2000). Controlling the intelligibility of referring expressions in dialogue. *Journal of Memory & Language*, *42*, 1-22. doi:10.1006/jmla.1999.2667
- BARD, E. G., & AYLETT, M. P. (1999). The disassociation of deaccenting, givenness, and syntactic role in spontaneous speech. In J. J. Ohala, Y. Hasegawa, M. Ohala, D. Granville, & A. C. Bailey (Eds.), *Proceedings of the XIV International Congress of Phonetic Sciences* (pp. 1753-1756). San Francisco.
- BELL, A., BRENIER, J. M., GREGORY, M. L., GIRAND, C., & JURAFSKY, D. (2009). Predictability effects on durations of content and function words in conversational English. *Journal of Memory & Language*, *60*, 92-111. doi:10.1016/j.jml.2008.06.003
- BELL, A., JURAFSKY, D., FOSLER-LUSSIER, E., GIRAND, C., GREGORY, M. L., & GILDEA, D. (2003). Effects of disfluencies, predictability,

- and utterance position on word form variation in English conversation. *Journal of the Acoustical Society of America*, **113**, 1001-1024. doi:10.1121/1.1534836
- BOCK, J. K., & LEVELT, W. (1994). Language production: Grammatical encoding. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 945-983). San Diego: Academic Press.
- BOERSMA, P., & WEENINK, D. (2007). Praat: Doing phonetics by computer (Version 4.5.14) [Computer program]. Retrieved February 5, 2007, from www.praat.org/.
- BROWN, P. M., & DELL, G. S. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology*, **19**, 441-472.
- DAHAN, D., TANENHAUS, M. K., & CHAMBERS, C. G. (2002). Accent and reference resolution in spoken-language comprehension. *Journal of Memory & Language*, **47**, 292-314. doi:10.1016/S0749-596X(02)00001-3
- DELL, G. S. (1990). Effects of frequency and vocabulary type on phonological speech errors. *Language & Cognitive Processes*, **5**, 313-349. doi:10.1080/01690969008407066
- FOSLER-LUSSIER, E., & MORGAN, N. (1999). Effects of speaking rate and word predictability on conversational pronunciations. *Speech Communication*, **29**, 137-158. doi:10.1016/S0167-6393(99)00035-7
- FOWLER, C. A. (1988). Differential shortening of repeated context words produced in various communicative contexts. *Language & Speech*, **31**, 307-319.
- FOWLER, C. A., & HOUSUM, J. (1987). Talkers' signaling of "new" and "old" words in speech and listeners' perception and use of the distinction. *Journal of Memory & Language*, **26**, 489-504. doi:10.1016/0749-596X(87)90136-7
- FRANK, A., & JAEGER, T. F. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *The 30th Annual Conference of the Cognitive Science Society* (pp. 933-938). Austin, TX: Cognitive Science Society.
- GREGORY, M. L., RAYMOND, W. D., BELL, A., FOSLER-LUSSIER, E., & JURAFSKY, D. (1999). The effects of collocational strength and contextual predictability in lexical production. In *Proceedings of the Chicago Linguistic Society* (Vol. 35, pp. 151-166). Chicago: University of Chicago Press.
- GRIFFIN, Z. M., & BOCK, J. K. (1998). Constraint, word frequency, and the relationship between lexical processing levels in spoken word production. *Journal of Memory & Language*, **38**, 313-338. doi:10.1006/jmla.1997.2547
- HORTON, W. S., & KEYSAR, B. (1996). When do speakers take into account common ground? *Cognition*, **59**, 91-117. doi:10.1016/0010-0277(96)81418-1
- ISAACS, A. M., & WATSON, D. G. (2009, March). *Speakers and listeners don't agree: Audience design in the production and comprehension of acoustic prominence*. Poster presentation at CUNY 2009: Conference on Human Sentence Processing, Davis, CA.
- ISAACS, A. M., & WATSON, D. G. (2010). Accent detection is a slippery slope: Direction and rate of  $f_0$  change drives listeners' comprehension. *Language & Cognitive Processes*, **25**, 1178-1200.
- ISAACS, E. A., & CLARK, H. H. (1987). References in conversations between experts and novices. *Journal of Experimental Psychology: General*, **116**, 26-37.
- JAEGER, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, **61**, 23-62. doi:10.1016/j.cogpsych.2010.02.002
- JESCHENIAK, J. D., & LEVELT, W. J. M. (1994). Word frequency effects in speech production: Retrieval of syntactic information and phonological form. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 824-843. doi:10.1037/0278-7393.20.4.824
- JURAFSKY, D., BELL, A., GREGORY, M., & RAYMOND, W. D. (2001). Probabilistic relations between words: Evidence from reduction in lexical production. In J. Bybee & P. Hopper (Eds.), *Frequency and the emergence of linguistic structure* (pp. 229-254). Amsterdam: Benjamins.
- KIDD, C., & JAEGER, T. F. (2008, April). Prosodic phrasing and function word pronunciation. Paper presented at Experimental and Theoretical Advances in Prosody, Cornell University, Ithaca, NY.
- LIEBERMAN, R. (1963). Some effects of the semantic and grammatical context on the production and perception of speech. *Language & Speech*, **6**, 172-175. doi:10.1121/1.1918465
- MARR, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- PIERREHUMBERT, J. (1980). *The phonology and phonetics of English intonation*. Unpublished doctoral dissertation.
- PIERREHUMBERT, J., & HIRSCHBERG, J. (1990). The meaning of intonational contours in the interpretation of discourse. In P. Cohen, J. Morgan, & M. Pollack (Eds.), *Intentions in communication* (pp. 271-311). Cambridge, MA: MIT Press.
- PLUYMAEKERS, M., ERNESTUS, M., & BAAYEN, R. H. (2005a). Articulatory planning is continuous and sensitive to informational redundancy. *Phonetica*, **62**, 146-159. doi:10.1159/000090095
- PLUYMAEKERS, M., ERNESTUS, M., & BAAYEN, R. H. (2005b). Lexical frequency and acoustic reduction in spoken Dutch. *Journal of the Acoustical Society of America*, **118**, 2561-2569. doi:10.1121/1.2011150
- ROSSIGN, B., & POURTOIS, G. (2001). Revisiting Snodgrass and Vanderwart's object database: Color and texture improve object recognition. *Journal of Vision*, **1**(3), 413a.
- SNODGRASS, J. G., & VANDERWART, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning & Memory*, **6**, 174-215. doi:10.1037/0278-7393.6.2.174
- WATSON, D. G. (2010). The many roads to prominence: Understanding emphasis in conversation. In B. Ross (Ed.), *The psychology of learning and motivation* (Vol. 52, pp. 163-183). Burlington: Academic Press.
- WATSON, D. G., ARNOLD, J. E., & TANENHAUS, M. K. (2008). Tic Tac TOE: Effects of predictability and importance on acoustic prominence in language production. *Cognition*, **106**, 1548-1557. doi:10.1016/j.cognition.2007.06.009
- WINGFIELD, A. (1968). Effect of frequency on identification and naming objects. *American Journal of Psychology*, **81**, 226-234. doi:10.2307/1421267
- ZIPE, G. K. (1929). Relative frequency as a determinant of phonetic change. *Harvard Studies in Classical Philology*, **15**, 1-95.

(Manuscript received September 26, 2009;  
revision accepted for publication April 30, 2010.)